Year One Progress Report for NASA Grant NNG06GH15G

Principal Investigator:
Robert J. Brunner
Department of Astronomy & NCSA
University of Illinois
1002 West Green Street
Urbana, IL 61853
(217) 244-6099

Summary:

This work is designed to explore the efficacy of alternative computational hardware to the challenges of cosmological data analysis and machine learning. During the first year, we have made remarkable progress. First, we have demonstrated an approximately 90x speed-up of a brute-force two-point angular correlation analysis by using an FPGA-based reconfigurable system. We also have started to port a k-nearest neighbor, a pixel-based power spectrum estimator, and a tree-based correlation algorithm to this hardware. In addition, we have begun to explore the use of a graphic processor unit (GPU) to the two-point correlation algorithm. Finally, we have demonstrated our novel work in several publications and presentations.

Review:

The first algorithm that we have studied is the angular two-point correlation function, which measures how strongly a dataset is clustered in angular coordinates relative to a random dataset. This algorithm has a long history in Cosmological analyses and is starting to see applications in other disciplines as well. To provide a benchmark, we first implemented a brute-force correlation estimator in C that was compiled and executed on a number of different architectures. Over the course of several months, PI Brunner and Co-Is Kindratenko and Myers met regularly to discuss different strategies to successfully port this algorithm to the SRC dual-MAP architecture. Specific issues that were include memory management, task-breakdown, distance binning, and floating-point precision. In the end, our most efficient implementation on the dual-MAP was a 42-bit, fixed-precision implementation that provided a nearly 90X speed-up over a single processor Xeon. Full details of this work are provided in the papers and presentations listed below.

We also implemented this algorithm on the SGI RASC RC100 architecture and were successful with a double-precision implementation that provided a XX speed-up, however, we were unable to successfully implement the fixed-precision approach. We are currently exploring the use of a GPU architecture to provide yet another implementation of this algorithm. We are also exploring ways to improve on our current results. First, we have only processed limited datasets (~100,000 sources) due to memory issues. As realistic datasets are currently exceeding this by at least an order-of-magnitude, we are exploring ways to more effectively manage data streams in these novel architectures. Second, a number of performance tricks (such as using a kd-tree) are generally applied in real-world analyses that produce significant performance improvements over the traditional brute-force approach. As a result, we are looking into including these within our HPRC efforts to produce more realistic performance comparisons.

A second algorithm that we have started to explore is the standard angular power spectrum, in particular that approach we are looking at currently is the pixel-based estimator as described in Huterer et al, 2001, ApJ, 555, 547. Graduate student Brett Hayes, who is funded under this proposal, has implemented this algorithm and tested it in several architectures, including Copper, one of NCSA's supercomputing resources. He also has attended a training session at SRC so that he can port this algorithm to the SRC MAP architecture.

Finally, we have started exploring additional algorithms to port, including a k-nearest neighbor classification algorithm, a KD-tree based correlation estimator, and an n-point correlation estimator.

Publications:

D. Meixner, V. Kindratenko, D. Pointer, On Using Simulink to Program SRC-6 Reconfigurable Computer, in Proc. Military and Aerospace Programmable Logic Device-MAPLD'06.

V. Kindratenko, A. Myers, and R. Brunner, Exploring coarse- and fine-grain parallelism on a high-performance reconfigurable computer, submitted to Parallel Computing, 2006.

V. Kindratenko, Mitrion-C Application Development on SGI Altix 350/RC100, submitted to IEEE Symposium on Field-Programmable Custom Computing Machines-FCCM'07.

Presentations:

R. Brunner, Addressing Cosmological Questions by using Reconfigurable Computing, Reconfigurable Systems Summer Institute, Urbana, IL, July 2006.

V. Kindratenko, A. Myers, R. Brunner, Poster: Exploring Coarse-grain and Fine-grain Parallelism on SRC-6 Reconfigurable Computer, Reconfigurable Systems Summer Institute, Urbana, IL, July 2006.

V. Kindratenko, Accelerating Scientific Applications with Reconfigurable Computing, SRC Users Meeting, Urbana, IL, July 2006.

V. Kindratenko, Advanced Computer Architectures for LSST: Reconfigurable Computing, Third Annual LSST All Hands Meeting, Menlo Park, CA, December 2006.